

Bayesian Optimal Approximate Message Passing to Recover Structured Sparse Signals

Martin Mayer, *Student Member, IEEE*, and Norbert Goertz, *Senior Member, IEEE*

Abstract—We present a novel compressed sensing recovery algorithm – termed Bayesian Optimal Structured Signal Approximate Message Passing (BOSSAMP) – that jointly exploits the prior distribution and the structured sparsity of a signal that shall be recovered from noisy linear measurements. Structured sparsity is inherent to *group sparse* and *jointly sparse* signals. Our algorithm is based on approximate message passing that poses a low complexity recovery algorithm whose Bayesian optimal version allows to specify a prior distribution for each signal component. We utilize this feature in order to establish an iteration-wise *extrinsic group update* step, in which likelihood ratios of neighboring group elements provide soft information about a specific group element. Doing so, the recovery of structured signals is drastically improved.

We derive the extrinsic group update step for a *sparse binary* and a *sparse Gaussian* signal prior, where the nonzero entries are either one or Gaussian distributed, respectively. We also explain how BOSSAMP is applicable to arbitrary sparse signals.

Simulations demonstrate that our approach exhibits superior performance compared to the current state of the art, while it retains a simple iterative implementation with low computational complexity.

Index Terms—compressed sensing, message passing, group sparse, jointly sparse, sparse binary, Bernoulli-Gaussian, Gaussian mixture, extrinsic information, turbo decoding

I. INTRODUCTION

Solving a linear system of equations $\mathbf{y} = \mathbf{A}\mathbf{x}$ for \mathbf{x} is an omnipresent problem in various fields, as a myriad of problem statements can be written in such form. While the very classical techniques such as least squares or Minimum Mean Squared Error (MMSE) estimation minimize the error with respect to the ℓ_2 -norm, the incorporation of additional knowledge like sparsity and structure of \mathbf{x} has gained a lot of attention over recent years.

In particular, compressed sensing was introduced in [1]–[3] to solve a linear system of equations in case of a sparsely populated \mathbf{x} , i.e., the vector features only few nonzero entries. A change of paradigm was triggered by recognizing that a sparse \mathbf{x} can be reconstructed perfectly from an *underdetermined* system of linear equations. In the context of signal processing, a signal vector \mathbf{x} can thus be reconstructed from undersampling, where the sampling basis functions are the columns of \mathbf{A} , and where the samples are stored in \mathbf{y} . This led to a huge popularity in the search for efficient recovery algorithms that incorporate the sparsity constraint; the classical compressed sensing formulation

$$\hat{\mathbf{x}} = \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^N} \|\tilde{\mathbf{x}}\|_0 \quad \text{s.t.} \quad \mathbf{A}\tilde{\mathbf{x}} = \mathbf{y},$$

where the pseudo norm $\|\tilde{\mathbf{x}}\|_0$ counts the number of nonzero entries in $\tilde{\mathbf{x}}$, leads to a combinatorial search for $\hat{\mathbf{x}}$ and is generally

NP-hard to solve. The problem was relaxed into an ℓ_1 -norm minimization called basis pursuit (denoising) [4] which is also applicable to noisy samples \mathbf{y} . An alternative formulation that introduces a controllable weight for the sparsity constraint is given by the Least Absolute Shrinkage and Selection Operator (LASSO) [5]. A computationally efficient recovery algorithm that iteratively solves the LASSO is provided by Approximate Message Passing (AMP) that was introduced in [6]–[9]. Its foundation is Gaussian loopy belief propagation [10] with simplified message passing that assumes high dimensional signal vectors \mathbf{x} . Its Bayesian optimal version that exploits the signal prior is described in [7]–[9], we will henceforth denote it as Bayesian optimal Approximate Message Passing (BAMP). BAMP was further extended in [11] in order to allow for non-linear and non-Gaussian output relations; the samples \mathbf{y} are transformed by an arbitrary but known function, and the output of the transformation is known to the estimator.

In many problems, \mathbf{x} features a certain structure in the sparsity, i.e., collections of entries (groups) contain either only zeros or only nonzero entries. Group sparsity typically occurs in (image) classification tasks [12]–[14]. To account for this, the LASSO was extended to the group LASSO [15]–[18]. Another established scheme is grouped orthogonal matching pursuit [19]. In the realm of AMP, the generalized approach [11] was extended in [20] to incorporate the signal structure.

Aside from group sparsity, structure is also found in jointly sparse signals, i.e., N_B vectors \mathbf{x}_b share a common support (the nonzero entries occur at the same indices, $\forall b \in \{1, \dots, N_B\}$). A prominent instance of this is the multiple measurement vector problem [21], [22]. Typical applications of the jointly sparse case are neuromagnetic imaging [21], [23] and direction-of-arrival estimation [24].

A. Contributions

In this paper, we present a novel recovery algorithm – termed Bayesian Optimal Structured Signal Approximate Message Passing (BOSSAMP) – that extends BAMP to incorporate the signal structure, such as group or joint sparsity. The approach also allows for overlapping groups, and it is not restricted to the sparse case. The key feature is the inclusion of a *group update* step that is inspired by an *extrinsic information* exchange that is predominantly used in coding [25]–[28], where it is also known as the *turbo principle*. In each iteration of BOSSAMP, the probability that a specific group entry was zero is updated by accumulating the extrinsic information of all other group entries in terms of likelihood ratios. This leads to a superior recovery performance compared to other

state of the art approaches such as [17]–[20], which we show by simulation. Furthermore, our algorithm converges to a solution in very few iterations, while its implementation stays simple and efficient. Specifically, it only requires two matrix-vector multiplications per iteration, matrix inversions are not required.

B. Related Work and Novelty

Merging (loopy) belief propagation with turbo equalization to recover structured sparse signals has already been suggested in [29], where the factor graph of the observation structure was extended by a pattern structure. Hidden binary indicators were used to model whether signal entries are active (nonzero) or inactive (zero). Sparsity pattern beliefs are exchanged between the observation structure and the sparsity structure in an iterative manner by leveraging (loopy) belief propagation. This approach was later utilized in compressive imaging [30] to exploit the sparsity and persistence across scales of 2D wavelet coefficients of natural images. A generalized manifestation of AMP that embeds the ideas of [29] and presents an algorithmic implementation is provided in [20], to which we compare our scheme to.

While [29], [30] approach the topic from the message passing point of view, we focus on an alternative description that utilizes likelihood ratios in terms of L -values. We utilize the classic BAMP algorithm and extend the iteration loop by two steps, namely the *group update* and the subsequent *prior update*. On the one hand, BAMP assumes independently but non-identically distributed signal entries in \mathbf{x} to perform scalar MMSE estimation. On the other hand, our two additional steps update the entry-wise prior information for the next BAMP iteration by exploiting the structure in the sparsity. In a “ping-pong” manner, an MMSE estimate emerges.

We give a detailed and easy-to-follow derivation of the group update step, specifically for the following two prominent cases, where we provide simple closed form expressions: first for the *sparse binary* case, where L -values can be formulated naturally, and then for the *sparse Gaussian* (also known as Bernoulli-Gaussian) case, where we introduce a latent binary variable that indicates whether an entry was zero or nonzero.

While the resulting BOSSAMP algorithm shares similarities with the Hybrid Generalized Approximate Message Passing (HGAMP) algorithm [20] which is applicable to a more general class of problems, our approach sticks to the standard AMP framework [6]–[9] which is mainly applicable to samples that are corrupted by Gaussian noise. While being restricted to a smaller class of problems, this leads to a *simpler implementation* and a *higher comprehensibility*. Furthermore, simulation results suggest that our approach outperforms HGAMP in terms of recovery performance and phase transitions.

C. Notation

Boldface letters such as \mathbf{A} and \mathbf{a} denote matrices and vectors, respectively. Considering matrix \mathbf{A} , $\mathbf{A}_{i,:}$ is its i -th row, while $\mathbf{A}_{:,j}$ is its j -th column. The superscript $(\cdot)^T$ denotes the transposition of a matrix or vector. The vectorization of

an $M \times N$ matrix is denoted $\mathbf{A}(\cdot) \equiv [\mathbf{A}_{:,1}^T, \dots, \mathbf{A}_{:,N}^T]^T$. The $N \times N$ identity matrix is denoted \mathbf{I}_N . The length N all-one vector is denoted $\mathbf{1}_N$, while the $N \times N$ all-one matrix is denoted $\mathbf{1}_{N \times N}$. Similarly, we define the all-zero vector $\mathbf{0}_N$ and the all-zero matrix $\mathbf{0}_{N \times N}$. Calligraphic letters \mathcal{S} denote sets, their usage as subscript $\mathbf{a}_{\mathcal{S}}$ implies that only the vector entries defined by the elements in \mathcal{S} are selected. The cardinality of a set is denoted by $|\mathcal{S}|$. Random variables and vectors are denoted by sans serif font as x and \mathbf{x} , respectively. While $x \sim \mathcal{N}(\mu, \sigma^2)$ denotes a Gaussian distributed random variable x with mean μ and variance σ^2 , the shorthand notation

$$\mathcal{N}(x|\mu, \sigma^2) \equiv \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x - \mu)^2\right) \quad (1)$$

denotes that such a Gaussian distribution is evaluated at the value x .

D. Outline

The remainder of this paper is outlined as follows: Section II reviews compressed sensing and draws the link between the probabilistic graphical model of the estimation and the iterative recovery schemes AMP and BAMP, the algorithmic implementation of these is also presented. Finally, two prominent sparse signal priors, namely the sparse binary and the sparse Gaussian prior, are introduced. Section III describes the group sparse case and presents the corresponding BOSSAMP algorithm, whose update rules are derived for the sparse binary and the sparse Gaussian signal prior, respectively. Section IV does the same for the jointly sparse case. Section V discusses the extension of BOSSAMP to the arbitrary signal case. Section VI introduces the figures of merit and the comparative schemes for simulation, and then presents and discusses the numerical results. The paper is concluded in Section VII.

II. RECOVERY OF SPARSE SIGNALS

A. Compressed Sensing

Compressed sensing was introduced in [1]–[3] to reconstruct a high-dimensional signal vector $\mathbf{x} \in \mathbb{R}^N$ from $M < N$ linear measurements

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}, \quad (2)$$

where $\mathbf{y} \in \mathbb{R}^M$ is the measurement vector, $\mathbf{A} \in \mathbb{R}^{M \times N}$ is the fixed sensing matrix and $\mathbf{w} \in \mathbb{R}^M$ is additive measurement noise. Signal vector \mathbf{x} is assumed to be K -sparse: it has at most K out of N nonzero entries, where $K \ll N$. This enables to recover \mathbf{x} from (2) although the system of equations is underdetermined. To ensure stable recovery from noisy measurements, the sensing matrix \mathbf{A} has to satisfy the Restricted Isometry Property (RIP) [3]

$$(1 - \delta_{2K}) \|\mathbf{v}\|_2^2 \leq \|\mathbf{A}\mathbf{v}\|_2^2 \leq (1 + \delta_{2K}) \|\mathbf{v}\|_2^2 \quad (3)$$

of order $2K$ and level $\delta_{2K} \in (0, 1)$ for all $2K$ -sparse vectors \mathbf{v} , which basically implies that the linear operator \mathbf{A} preserves the Euclidean distance between every pair of K -sparse vectors up to a small constant δ_{2K} .

An appropriate sensing matrix can be constructed by picking \mathbf{A} randomly with i.i.d. (sub-)Gaussian entries. Such matrices were proven in [31] to almost surely satisfy the RIP while the number of required measurements for successful recovery is lower bounded by

$$M = \left\lceil cK \log \frac{N}{K} \right\rceil, \quad (4)$$

where c is a small constant and the ceiling operation $\lceil \cdot \rceil$ ensures an integer number of measurements.

B. Graphical Model

A graphical model poses the probabilistic foundation of compressed sensing recovery algorithms that are based on message passing, such as AMP and BAMP, see [32]. The key ingredient is a factorization of a multivariate distribution – in our case the posterior distribution of \mathbf{x} given \mathbf{y} based on (2) – into many factors. This factorization is described and visualized by a factor graph [33], [34], which we will now present.

Let us begin with the underlying probabilistic assumptions. Considering (2), we assume that only sensing matrix \mathbf{A} is deterministic (fixed), which leaves us to characterize the distributions of measurement vector $\mathbf{y} = [y_1, \dots, y_m, \dots, y_M]^T$, signal vector $\mathbf{x} = [x_1, \dots, x_n, \dots, x_N]^T$ and noise vector $\mathbf{w} = [w_1, \dots, w_m, \dots, w_M]^T$.

We assume white Gaussian noise with zero mean and covariance $\sigma_w^2 \mathbf{I}_M$, i.e., $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}_M)$. The noise Probability Density Function (PDF) calculates as

$$f_{\mathbf{w}}(\mathbf{w}) = \prod_{m=1}^M f_w(w_m) = \prod_{m=1}^M \mathcal{N}(w_m | 0, \sigma_w^2). \quad (5)$$

The joint PDF of signal and measurement can be factored according to Bayes' rule:

$$f_{\mathbf{x}, \mathbf{y}}(\mathbf{x}, \mathbf{y}) = \underbrace{f_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y})}_{\text{posterior}} f_{\mathbf{y}}(\mathbf{y}) = \underbrace{f_{\mathbf{y}|\mathbf{x}}(\mathbf{y}|\mathbf{x})}_{\text{likelihood}} \underbrace{f_{\mathbf{x}}(\mathbf{x})}_{\text{prior}}. \quad (6)$$

We assume independently distributed signal entries with PDF $f_{x_n}(x_n)$, the *prior*, thus, factors as

$$f_{\mathbf{x}}(\mathbf{x}) = \prod_{n=1}^N f_{x_n}(x_n). \quad (7)$$

The *likelihood* is characterized by the noise PDF:

$$f_{\mathbf{y}|\mathbf{x}}(\mathbf{y}|\mathbf{x}) = f_{\mathbf{w}}(\mathbf{y} - \mathbf{A}\mathbf{x}) = \prod_{m=1}^M f_w(y_m - \mathbf{A}_{m,:}\mathbf{x}). \quad (8)$$

Estimators of \mathbf{x} typically rely on the *posterior*

$$f_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y}) = \prod_{m=1}^M \frac{1}{f_{y_m}(y_m)} f_{y_m|\mathbf{x}}(y_m|\mathbf{x}) f_{\mathbf{x}}(\mathbf{x}) \quad (9)$$

that entails M factors.

The resulting factor graph $\text{FG} = (\mathcal{V}, \mathcal{F}, \mathcal{E})$ consists of the variable nodes $\mathcal{V} = \{1, \dots, N\}$ that encompass the signal of interest \mathbf{x} , the factor nodes $\mathcal{F} = \{1, \dots, M\}$ associated to (9),




Fig. 1. Factor graph of measurement (2) associated to (9).

and the edges $\mathcal{E} = \{(m, n) : m \in \mathcal{F}, n \in \mathcal{V}\}$ that correspond to the relations between \mathbf{x} and \mathbf{y} which are dictated by the entries of \mathbf{A} , i.e., $\mathbf{A}_{m,n}$. Since typically, all entries of \mathbf{A} are nonzero, our bipartite graph, as illustrated by Figure 1, is fully connected.

As we intend to recover \mathbf{x} from \mathbf{y} knowing \mathbf{A} , we formulate the MMSE estimator

$$\hat{\mathbf{x}}_{\text{MMSE}}(\mathbf{y}) = \mathbb{E}_{\mathbf{x}}[\mathbf{x}|\mathbf{y} = \mathbf{y}] = \int_{\mathbb{R}^N} \tilde{\mathbf{x}} f_{\mathbf{x}|\mathbf{y}}(\tilde{\mathbf{x}}|\mathbf{y}) d\tilde{\mathbf{x}}. \quad (10)$$

This task can be approximately¹ solved utilizing message passing (belief propagation) and the sum-product algorithm [33]–[35] on our factor graph. However, a computationally efficient method is only obtained after a series of assumptions and approximations – described in [9] – that yield the AMP algorithm. Note that AMP performs *scalar* MMSE estimation independently for every component of \mathbf{x} .

C. Approximate Message Passing (AMP)

AMP has been introduced in [6]–[9] to efficiently solve the LASSO problem [5], also known as basis pursuit denoising [4], that constitutes a non-linear convex optimization problem

$$\hat{\mathbf{x}}_{\text{LASSO}}(\mathbf{y}; \lambda) = \arg \min_{\tilde{\mathbf{x}}} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{A}\tilde{\mathbf{x}}\|_2^2 + \lambda \|\tilde{\mathbf{x}}\|_1 \right\}. \quad (11)$$

The underlying intuition is to find the most accurate solution with the smallest support (motivated by the assumed sparsity of \mathbf{x}), where λ allows for a trade-off between accuracy with respect to the ℓ_2 observation error $\|\mathbf{y} - \mathbf{A}\tilde{\mathbf{x}}\|_2^2$ and the sparsity² $\|\tilde{\mathbf{x}}\|_1$ of the solution.

An illustrative approach to obtain the LASSO from MMSE estimator (10) is to assume that the signal vector entries are Laplacian distributed, i.e., $f_{x_n}(x_n) = f_{\text{Laplace}}(x_n; 0, \kappa)$ with

$$f_{\text{Laplace}}(x; \mu, \kappa) = \frac{1}{2\kappa} \exp\left(-\frac{1}{\kappa}|x - \mu|\right). \quad (12)$$

The zero mean Laplace distribution poses a sparsity enforcing prior, i.e., its probability mass is concentrated around zero. Plugging the Laplace signal prior into (9) and calculating the MMSE estimate (10), we obtain the LASSO (11) with $\lambda = \sigma_w^2/\kappa$.

¹The considered graph typically contains cycles which lead to loopy belief propagation that yields an approximate result.

²Sparsity is usually expressed by the ℓ_0 -“norm” according to $\|\mathbf{x}\|_0 \leq K$. The ℓ_1 -norm relaxation was proven in [1] to very often yield the same result in high dimensions, while introducing a favorable *convex* optimization problem.

In AMP, λ is a design parameter. For the optimal choice of λ , it was shown in [36] that the fixed point of the AMP solution coincides with the LASSO solution in the *asymptotic regime* where $M/N = \text{const.}$ while $N, M \rightarrow \infty$.

As discussed in [32], [37], AMP *decouples* the estimation problem associated to measurement (2) into N uncoupled scalar problems in the asymptotic regime:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \xrightarrow{\text{asympt.}} \begin{cases} u_1 = x_1 + \tilde{w}_1 \\ \vdots \\ u_N = x_N + \tilde{w}_N \end{cases}, \quad (13)$$

where the *effective noise* asymptotically obeys $\tilde{w}_n \sim \mathcal{N}(0, \beta)$. Note that it is Gaussian³, and $\beta > \sigma_w^2$. Revisiting the LASSO problem in the scalar case, i.e.,

$$\hat{x}_{\text{LASSO}}(u; \lambda) = \arg \min_{\tilde{x}} \left\{ \frac{1}{2} (u - \tilde{x})^2 + \lambda |\tilde{x}| \right\}, \quad (14)$$

it is known that the *soft thresholding* function

$$\eta(u; \tau) = \begin{cases} u + \tau & \text{if } u < -\tau \\ 0 & \text{if } -\tau \leq u \leq \tau \\ u - \tau & \text{if } u > \tau \end{cases} \quad (15)$$

admits a (possibly optimal) solution to (14), see [5]. The soft thresholding function acts as a *denoiser*, i.e., it sets values below a certain threshold to zero. This function is also found in the AMP algorithm – our implementation is stated by Algorithm 1 – where it is applied entry-wise on the decoupled measurements $\mathbf{x} + \mathbf{A}^T \mathbf{r}$, i.e., on $u_n = x_n + (\mathbf{A}_{:,n})^T \mathbf{r}$. The iterations are stopped once the change in the estimated signal is below a certain threshold – controlled by ϵ_{tol} – or the maximal number of iterations t_{max} is reached. For a detailed derivation of AMP, the interested reader is referred to [6]–[9].

D. Bayesian optimal Approximate Message Passing (BAMP)

While the standard AMP algorithm implicitly assumes a (sparsity enforcing) Laplacian signal prior, its Bayesian optimal version BAMP allows to specify arbitrary signal priors $f_{x_n}(x_n)$, individually for each entry of the signal vector — this is a key feature that will be exploited by the proposed algorithms for structured sparsity. As before, we will stick to the main features and refer to [7]–[9] for details.

BAMP builds on the same decoupling principle [32], [37] as AMP, which is valid in the asymptotic regime and approximately satisfied in finite dimensions. Let us discuss the decoupled scalar problem (13), $u_n = x_n + \tilde{w}_n$, where we know that $\tilde{w}_n \sim \mathcal{N}(0, \beta)$. The (entry-wise) *posterior* of the decoupled problem thus reads, using Bayes' theorem,

$$f_{x_n|u_n}(x_n|u_n) = \frac{1}{f_{u_n}(u_n)} f_{u_n|x_n}(u_n|x_n) f_{x_n}(x_n), \quad (16)$$

where $f_{u_n}(u_n) = \int_{-\infty}^{\infty} f_{u_n|x_n}(u_n|\tilde{x}_n) f_{x_n}(\tilde{x}_n) d\tilde{x}_n$ and $f_{u_n|x_n}(u_n|x_n) = \mathcal{N}(u_n|x_n, \beta)$. Instead of soft thresholding,

³This is an assumption that is satisfied in the asymptotic regime.

Algorithm 1 AMP

```

1: initialize  $\mathbf{x}^t = \mathbf{0}$  and  $\mathbf{r}^t = \mathbf{y}$  for  $t = 0$ 
2: do
3:    $t = t + 1$  ▷ advance iterations
4:    $\tau = \frac{\lambda}{\sqrt{M}} \|\mathbf{r}^{t-1}\|_2$  ▷ compute threshold
5:    $\mathbf{x}^t = \eta(\mathbf{x}^{t-1} + \mathbf{A}^T \mathbf{r}^{t-1}; \tau)$  ▷ soft thresholding
6:    $b = \frac{1}{M} \|\mathbf{x}^t\|_0$  ▷ compute sparsity
7:    $\mathbf{r}^t = \mathbf{y} - \mathbf{A}\mathbf{x}^t + b\mathbf{r}^{t-1}$  ▷ compute residual
8: while  $\|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2 > \epsilon_{\text{tol}} \|\mathbf{x}^{t-1}\|_2$  and  $t < t_{\text{max}}$ 
9: return  $\hat{\mathbf{x}} = \mathbf{x}^t$  ▷ recovered sparse vector

```

Algorithm 2 BAMP

```

1: initialize  $\mathbf{x}^t = \mathbf{0}$  and  $\mathbf{r}^t = \mathbf{y}$  for  $t = 0$ 
2: do
3:    $t = t + 1$ 
4:    $\mathbf{u}^{t-1} = \mathbf{x}^{t-1} + \mathbf{A}^T \mathbf{r}^{t-1}$  ▷ decoupled measurements
5:    $\beta^{t-1} = \frac{1}{M} \|\mathbf{r}^{t-1}\|_2^2$  ▷ effective noise estimate
6:    $\mathbf{x}^t = F(\mathbf{u}^{t-1}; \beta^{t-1})$  ▷ estimate signal
7:    $\mathbf{r}^t = \mathbf{y} - \mathbf{A}\mathbf{x}^t + \mathbf{r}^{t-1} \frac{1}{M} \sum_{n=1}^N F'(u_n^{t-1}; \beta^{t-1})$ 
8: while  $\|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2 > \epsilon_{\text{tol}} \|\mathbf{x}^{t-1}\|_2$  and  $t < t_{\text{max}}$ 
9: return  $\hat{\mathbf{x}} = \mathbf{x}^t$ 

```

BAMP utilizes the following functions [7]–[9]:

$$F(u_n; \beta) = \mathbb{E}_{x_n} \{x_n | u_n = u_n; \beta\}, \quad (17)$$

$$G(u_n; \beta) = \text{Var}_{x_n} \{x_n | u_n = u_n; \beta\}, \quad (18)$$

$$F'(u_n; \beta) = \frac{d}{du_n} F(u_n; \beta). \quad (19)$$

The conditional expectation (17) yields the scalar MMSE estimate of x_n given the decoupled measurement u_n . Our implementation of BAMP is stated by Algorithm 2. Function (17) is applied entry-wise on a vector argument, (18) is typically an intermediate step to compute (19). Note that for a specified signal prior, functions (17)–(19) admit closed form expressions. We will now specify those for a *sparse binary* and a *sparse Gaussian* signal prior, respectively.

Sparse Binary Signal Prior: for $x_n \in \{0, 1\}$, the prior reads

$$f_{x_n}(x_n) = \gamma_n \delta(x_n) + (1 - \gamma_n) \delta(x_n - 1), \quad (20)$$

where γ_n indicates the probability of x_n being a zero. In this setting, functions (17)–(19) boil down to

$$F(u_n; \beta, \gamma_n) = \frac{1}{1 + \exp\left(\frac{1-2u_n}{2\beta} + \log \frac{\gamma_n}{1-\gamma_n}\right)}, \quad (21)$$

$$G(u_n; \beta, \gamma_n) = F(u_n; \beta, \gamma_n) - F(u_n; \beta, \gamma_n)^2, \quad (22)$$

$$F'(u_n; \beta, \gamma_n) = \frac{1}{\beta} G(u_n; \beta, \gamma_n). \quad (23)$$

Sparse Gaussian Signal Prior: for $x_n \in \{0, \mathcal{N}(0, \sigma_{x_n}^2)\}$, the prior reads

$$f_{x_n}(x_n) = \gamma_n \delta(x_n) + (1 - \gamma_n) \mathcal{N}(x_n | 0, \sigma_{x_n}^2), \quad (24)$$

where γ_n indicates the probability of x_n being a zero. In literature, this case is also known as the *Bernoulli-Gaussian* case. Functions (17)–(19) calculate as

$$F(u_n; \beta, \gamma_n) = u_n M(u_n, \gamma_n, q), \quad (25)$$

$$G(u_n; \beta, \gamma_n) = \beta M(u_n, \gamma_n, q) + m(u_n, \gamma_n, q), \quad (26)$$

$$F'(u_n; \beta, \gamma_n) = \frac{1}{\beta} G(u_n; \beta, \gamma_n), \quad (27)$$

with $q = \frac{\sigma_{x_n}^2}{\beta}$ and the auxiliary functions

$$M(u_n, \gamma_n, q) = \frac{q}{1+q} \frac{1}{1+m(u_n, \gamma_n, q)}, \quad (28)$$

$$m(u_n, \gamma_n, q) = \frac{\gamma_n}{1-\gamma_n} \sqrt{1+q} \exp\left(-\frac{u_n^2}{2\beta} \frac{q}{1+q}\right). \quad (29)$$

If the number of nonzero entries K is known a priori, we choose

$$\gamma_n = 1 - \frac{K}{N}, \forall n \in \mathcal{V} = \{1, \dots, N\}. \quad (30)$$

In case of unknown sparsity, one has to assume a certain sparsity and plug in an estimate for K .

III. RECOVERY OF GROUP SPARSE SIGNALS

In the group sparse case, signal vector \mathbf{x} is partitioned into N_G groups such that the groups partition (non-overlapping case) the total support set $\mathcal{V} = \{1, \dots, N\}$:

$$\mathcal{V} = \bigcup_{g=1}^{N_G} \mathcal{G}_g, \quad (31)$$

where \mathcal{G}_g contains the signal vector indices that correspond to group g . In case of overlapping groups, the intersection of two different groups may contain elements. The signal vector entries that correspond to a group are either *all zero* or *all nonzero* — knowing the groups reflects the a priori knowledge of the *signal structure*. An exemplary group assignment on a factor graph is depicted in Figure 2. While the groups may vary in their size $|\mathcal{G}_g|$, the signal vector \mathbf{x} is assumed to be K -sparse, which typically implies that the number of nonzero groups is small.

Considering BAMP in the sparse signal case, the priors (20) and (24) allow to modify the probability that a signal entry x_n is zero, for each entry of \mathbf{x} individually. This is the key feature exploited by Bayesian Optimal Structured Signal Approximate Message Passing (BOSSAMP).

A. BOSSAMP and Group Sparse Binary Signals

Binary signals allow for a convenient computation of soft information in terms of L -values⁴ [25], [26]:

$$L(x_n) = \log \frac{P(x_n = 0)}{P(x_n = 1)} = \log \frac{\gamma_n}{1-\gamma_n}. \quad (32)$$

A large positive value indicates a high probability of x_n being a zero, a large negative value a high probability of x_n being

⁴ L -values are log likelihood ratios in the context of coding. They are typically used in soft-input channel decoding or in iterative decoding.




Fig. 2. Factor graph with group indication.

a one. If we consider the decoupled measurements (13), the conditional L -values read, using (21),

$$\begin{aligned} L(x_n | u_n = u_n) &= \log \frac{P(x_n = 0 | u_n = u_n)}{P(x_n = 1 | u_n = u_n)} \\ &= \log \frac{1 - F(u_n; \beta, \gamma_n)}{F(u_n; \beta, \gamma_n)} \\ &= \frac{1 - 2u_n}{2\beta} + \log \frac{\gamma_n}{1 - \gamma_n}. \end{aligned} \quad (33)$$

If we compute the conditional L -values in each iteration of BAMP, i.e., compute the likelihood of being a zero for each signal entry estimate, we obtain for entry x_n at iteration t (cf. line 6 of Algorithm 2, using (33)):

$$L(x_n | u_n = u_n^{t-1}) = \frac{1 - 2u_n^{t-1}}{2\beta^{t-1}} + \log \frac{\gamma_n^{t-1}}{1 - \gamma_n^{t-1}}. \quad (34)$$

The key feature of BOSSAMP is to use the L -values (34) of iteration t as *extrinsic a priori information* [26] in the subsequent iteration $t + 1$. To that end, we calculate the L -values that accommodate the *innovation* of the new iteration:

$$L_n^t = L(x_n | u_n = u_n^{t-1}) - \underbrace{\log \frac{\gamma_n^{t-1}}{1 - \gamma_n^{t-1}}}_{\bar{L}_n^{t-1}} = \frac{1 - 2u_n^{t-1}}{2\beta^{t-1}}. \quad (35)$$

To exploit the group structure, we introduce the **binary extrinsic group update**

$$\begin{aligned} \bar{L}_n^t &= U_G(\mathbf{u}^{t-1}, \beta^{t-1}, \gamma_n^0) := \bar{L}_n^0 + \sum_{l \in \mathcal{G}_g \setminus n} L_l^t \\ &= \log \frac{\gamma_n^0}{1 - \gamma_n^0} + \sum_{l \in \mathcal{G}_g \setminus n} \frac{1 - 2u_l^{t-1}}{2\beta^{t-1}}, \end{aligned} \quad (36)$$

$$\forall n \in \mathcal{G}_g, \forall g \in \{1, \dots, N_G\},$$

which yields $\bar{\mathbf{L}}^t = [\bar{L}_1^t, \dots, \bar{L}_N^t]^T$. This can be interpreted as follows: \bar{L}_n^0 is the static prior knowledge about the n -th entry, and $\sum_{l \in \mathcal{G}_g \setminus n} L_l^t$ is the extrinsic information of the rest of the group that contains the innovation of the current iteration. If the extrinsic information provides a positive L -value, entry x_n becomes more likely to be a zero rather than a one.

After the extrinsic group update, the signal prior is updated accordingly for the subsequent iteration. We therefore introduce the **prior update**

$$\gamma_n^t = U_P(\bar{L}_n^t) := \frac{1}{1 + \exp(-\bar{L}_n^t)}, \forall n \in \mathcal{V}, \quad (37)$$

Algorithm 3 BOSSAMP for Group Sparse Signals

```

1: init.  $\mathbf{x}^t = \mathbf{0}_N$ ,  $\mathbf{r}^t = \mathbf{y}$  and  $\gamma^t = \mathbf{1}_N - \frac{K}{N}$  for  $t = 0$ 
2: do
3:    $t = t + 1$ 
4:    $\mathbf{u}^{t-1} = \mathbf{x}^{t-1} + \mathbf{A}^T \mathbf{r}^{t-1}$ 
5:    $\beta^{t-1} = \frac{1}{M} \|\mathbf{r}^{t-1}\|_2^2$ 
6:    $\mathbf{x}^t = F(\mathbf{u}^{t-1}; \beta^{t-1}, \gamma^{t-1})$ 
7:    $\mathbf{r}^t = \mathbf{y} - \mathbf{A} \mathbf{x}^t + \mathbf{r}^{t-1} \frac{1}{M} \sum_{n=1}^N F'(u_n^{t-1}; \beta^{t-1}, \gamma_n^{t-1})$ 
8:    $\bar{\mathbf{L}}^t = U_G(\mathbf{u}^{t-1}, \beta^{t-1}, \gamma^0) \quad \triangleright$  extrinsic group update
9:    $\gamma^t = U_P(\bar{\mathbf{L}}^t) \quad \triangleright$  prior update
10: while  $\|\mathbf{x}^t - \mathbf{x}^{t-1}\|_2 > \epsilon_{\text{tol}} \|\mathbf{x}^{t-1}\|_2$  and  $t < t_{\text{max}}$ 
11: return  $\hat{\mathbf{x}} = \mathbf{x}^t$ 

```

where we used L -value definition (32).

By including these two steps in the BAMP algorithm, we obtain the BOSSAMP algorithm for group sparse signals that is outlined in Algorithm 3 — functions (21), (23) and (36) are utilized for sparse binary signals. The zero probabilities are initialized as $\gamma_n^0 = 1 - \frac{K}{N}$, $\forall n \in \mathcal{V}$, according to (30).

B. BOSSAMP and Group Sparse Gaussian Signals

In the sparse Gaussian case, we are not able to express L -values directly as in the binary case. We tackle this problem by introducing a binary latent random variable inspired by the E-step of the Expectation Maximization (EM) algorithm [38]–[40]. This allows us to estimate the zero probabilities γ_n^t in each iteration t of BOSSAMP. Consider the prior distribution of the decoupled measurements $u_n = x_n + \tilde{w}_n$ (13) which can be expressed as a *Gaussian mixture* [39], [40]:

$$\begin{aligned} f_{u_n}(u_n) &= \gamma_n \mathcal{N}(u_n|0, \beta) + (1 - \gamma_n) \mathcal{N}(u_n|0, \beta + \sigma_{x_n}^2) \\ &= \sum_{i=1}^2 \alpha_{n,i} \mathcal{N}(u_n|\mu_i, \sigma_i^2). \end{aligned} \quad (38)$$

We distinguish between two Gaussian distributions:

- Distribution $i = 1$ is associated to the *zero entries* in the original signal \mathbf{x} . The corresponding estimates u_n solely contain the effective noise $\tilde{w}_n \sim \mathcal{N}(0, \beta)$ (i.e., $x_n = 0$). We thus set $\mu_1 = 0$ and $\sigma_1^2 = \beta$.
- Distribution $i = 2$ is associated to the *nonzero entries* where u_n contains the noisy signal entries ($x_n \neq 0$). Therefore, $\mu_2 = 0$ and $\sigma_2^2 = \beta + \sigma_{x_n}^2$.

The mixing coefficients $\alpha_{n,i}$ determine the probability of the individual mixture components: $\alpha_{n,1}$ is the probability of a zero entry, and $\alpha_{n,2} = 1 - \alpha_{n,1}$ is the probability of a nonzero entry. In order to estimate these probabilities, a *latent binary random variable* $z_{n,i} \in \{0, 1\}$, $i = 1, 2$, is introduced:

$$f_{u_n}(u_n) = \sum_{z_{n,i}} p_{z_{n,i}}(z_{n,i}) f_{u_n|z_{n,i}}(u_n|z_{n,i}), \quad (39)$$

where $z_{n,1} + z_{n,2} = 1$ and $f_{u_n|z_{n,i}}(u_n|z_{n,i} = 1)$ is a Gaussian distribution with mean μ_i and variance σ_i^2 . Defining the Probability Mass Function (PMF) of $z_{n,i}$ as

$$p_{z_{n,i}}(z_{n,i}) = (1 - \alpha_{n,i}) \delta(z_{n,i}) + \alpha_{n,i} \delta(z_{n,i} - 1), \quad (40)$$

the marginalization (39) becomes equivalent to (38) — we successfully reformulated (see [39]) the Gaussian mixture to involve a binary latent variable that can be estimated by the E-step of the EM algorithm. The E-step computes the probabilities

$$\begin{aligned} P(z_{n,i} = 1 | u_n = u_n) &= \frac{p_{z_{n,i}}(1) f_{u_n|z_{n,i}}(u_n | z_{n,i} = 1)}{f_{u_n}(u_n)} \\ &= \frac{\alpha_{n,i} \mathcal{N}(u_n | \mu_i, \sigma_i^2)}{\sum_{j=1}^2 \alpha_{n,j} \mathcal{N}(u_n | \mu_j, \sigma_j^2)}, \end{aligned} \quad (41)$$

which are called *responsibilities*; $P(z_{n,i} = 1 | u_n = u_n)$ is the responsibility of mixture component i for explaining observation u_n . They are used to estimate the zero probabilities γ_n^t in each iteration t , which is our *E-step*. The estimate reads

$$\begin{aligned} \gamma_n^t &:= P(z_{n,1} = 1 | u_n = u_n^{t-1}; \alpha_{n,1} = \gamma_n^{t-1}) \\ &= \frac{\gamma_n^{t-1} \mathcal{N}(u_n^{t-1} | 0, \beta^{t-1})}{\gamma_n^{t-1} \mathcal{N}(u_n^{t-1} | 0, \beta^{t-1}) + (1 - \gamma_n^{t-1}) \mathcal{N}(u_n^{t-1} | 0, \sigma_{x_n}^2 + \beta^{t-1})}. \end{aligned} \quad (42)$$

As our latent variable $z_{n,1}$ is binary, we can formulate the L -values

$$\begin{aligned} L(z_{n,1} | u_n = u_n^{t-1}) &= \log \frac{P(z_{n,1} = 1 | u_n = u_n^{t-1})}{P(z_{n,1} = 0 | u_n = u_n^{t-1})} \\ &= \log \frac{\tilde{\gamma}_n^t}{1 - \tilde{\gamma}_n^t} \end{aligned} \quad (43)$$

that indicate how likely signal entry x_n was to be zero (implies $z_{n,1} = 1$) given the measurement u_n^{t-1} . Similar to (35), we introduce the innovation L -values

$$\begin{aligned} L_n^t &= L(z_{n,1} | u_n = u_n^{t-1}) - \log \frac{\gamma_n^{t-1}}{1 - \gamma_n^{t-1}} \\ &= \log \frac{\tilde{\gamma}_n^t (1 - \gamma_n^{t-1})}{\gamma_n^{t-1} (1 - \tilde{\gamma}_n^t)} \\ &= \log \frac{\mathcal{N}(u_n^{t-1} | 0, \beta^{t-1})}{\mathcal{N}(u_n^{t-1} | 0, \sigma_{x_n}^2 + \beta^{t-1})} \\ &= \frac{1}{2} \log \frac{\beta^{t-1} + \sigma_{x_n}^2}{\beta^{t-1}} - \frac{1}{2} \frac{(u_n^{t-1})^2 \sigma_{x_n}^2}{\beta^{t-1} (\beta^{t-1} + \sigma_{x_n}^2)}. \end{aligned} \quad (44)$$

They are utilized for the **Gaussian extrinsic group update**

$$\begin{aligned} \bar{\mathbf{L}}_n^t &= U_G(\mathbf{u}^{t-1}, \beta^{t-1}, \gamma_n^0) := \bar{\mathbf{L}}_n^0 + \sum_{l \in \mathcal{G}_g \setminus n} L_l^t \\ &= \log \frac{\gamma_n^0}{1 - \gamma_n^0} + \frac{1}{2} \sum_{l \in \mathcal{G}_g \setminus n} \log \frac{\beta^{t-1} + \sigma_{x_l}^2}{\beta^{t-1}} - \frac{(u_l^{t-1})^2 \sigma_{x_l}^2}{\beta^{t-1} (\beta^{t-1} + \sigma_{x_l}^2)}, \\ &\quad \forall n \in \mathcal{G}_g, \forall g \in \{1, \dots, N_G\}, \end{aligned} \quad (45)$$

which is similar to (36). The subsequent prior update is the same as in the binary case, see (37).

The same BOSSAMP algorithm body as stated by Algorithm 3 is used — note that the functions (25), (27) and (45) are utilized for sparse Gaussian signals.

Algorithm 4 BOSSAMP for Jointly Sparse Signals

```

1: init.  $\mathbf{X}^t = \mathbf{0}_{N \times N_B}$ ,  $\mathbf{\Gamma}^t = \mathbf{1}_{N \times N_B} - \frac{K}{N}$  and  $\mathbf{r}_b^t = \mathbf{y}_b$ 
    $\forall b \in \mathcal{B} = \{1, \dots, N_B\}$  and  $t = 0$ 
2: do
3:    $t = t + 1$ 
4:   for  $b = 1$  to  $N_B$  do  $\triangleright$  BAMP iteration for all blocks
5:      $\mathbf{u}_b^{t-1} = \mathbf{x}_b^{t-1} + \mathbf{A}^{(b)T} \mathbf{r}_b^{t-1}$ 
6:      $\beta_b^{t-1} = \frac{1}{M} \|\mathbf{r}_b^{t-1}\|_2^2$ 
7:      $\mathbf{x}_b^t = F(\mathbf{u}_b^{t-1}; \beta_b^{t-1}, \gamma_b^{t-1})$ 
8:      $\mathbf{r}_b^t = \mathbf{y}_b - \mathbf{A}^{(b)} \mathbf{x}_b^t + \mathbf{r}_b^{t-1} \frac{1}{M} \sum_n F'(u_{n,b}^{t-1}; \beta_b^{t-1}, \gamma_b^{t-1})$ 
9:    $\bar{\mathbf{L}}^t = U_G(\mathbf{U}^{t-1}, \boldsymbol{\beta}^{t-1}, \mathbf{\Gamma}^0)$   $\triangleright$  extrinsic group update
10:   $\mathbf{\Gamma}^t = U_P(\bar{\mathbf{L}}^t)$   $\triangleright$  prior update
11: while  $\|\mathbf{X}(:,)^t - \mathbf{X}(:,)^{t-1}\|_2 > \epsilon_{\text{tol}} \|\mathbf{X}(:,)^{t-1}\|_2$  and  $t < t_{\text{max}}$ 
12: return  $\hat{\mathbf{x}} = \mathbf{x}^t$ 

```




Fig. 3. Factor graph of two jointly sparse signal vectors \mathbf{x}_1 and \mathbf{x}_2 .

IV. RECOVERY OF JOINTLY SPARSE SIGNALS

In the jointly sparse case, we consider N_B signal vectors $\mathbf{x}_b \in \mathbb{R}^N, b \in \mathcal{B} = \{1, \dots, N_B\}$, that share a *common support*

$$\mathcal{S}_{\mathbf{x}} \equiv \mathcal{S}_{\mathbf{x}_b}, \forall b \in \mathcal{B}, \quad (46)$$

where $\mathcal{S}_{\mathbf{x}_b}$ contains the indices of the nonzero entries in \mathbf{x}_b . In the most general case, the compressive measurements are formulated similar to (2) as

$$\mathbf{y}_b = \mathbf{A}^{(b)} \mathbf{x}_b + \mathbf{w}_b, \quad (47)$$

where $\mathbf{y}_b \in \mathbb{R}^M$, $\mathbf{A}^{(b)} \in \mathbb{R}^{M \times N}$, and $\mathbf{w}_b \in \mathbb{R}^M$. Let us collect the data blocks in matrices: $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_b, \dots, \mathbf{y}_{N_B}]$, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_b, \dots, \mathbf{x}_{N_B}]$ and $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_b, \dots, \mathbf{w}_{N_B}]$. If all *sensing matrices are equal*, i.e., $\mathbf{A} \equiv \mathbf{A}^{(b)}, \forall b \in \mathcal{B}$, we can rewrite (47) as

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{W}. \quad (48)$$

The joint sparsity is expressed by the *rows* of matrix \mathbf{X} , whose entries are either *all zero* or *all nonzero*. These rows can be interpreted as $N_G = N$ groups, where each group $\mathcal{G}_g, g \in N_G$, contains $|\mathcal{B}| = N_B$ elements. The BOSSAMP algorithm that exploits the joint sparsity is thus very similar to the one in the group sparse case. An exemplary factor graph with $N_B = 2$ blocks is depicted in Figure 3.

A. BOSSAMP and Jointly Sparse Binary Signals

Compared to the group sparse case, we essentially have to extend the indexing from vectors to matrices. In particular, $\mathbf{\Gamma}^t = [\gamma_1^t, \dots, \gamma_{N_B}^t]$ contains the zero probabilities of iteration t , where $(\mathbf{\Gamma}^t)_{n,b} = \gamma_{n,b}^t$ is the n -th entry of block b . Similarly, $(\mathbf{U}^t)_{n,b} = u_{n,b}^t$ and $(\bar{\mathbf{L}}^t)_{n,b} = \bar{L}_{n,b}^t$, and $\boldsymbol{\beta}^t = [\beta_1^t, \dots, \beta_{N_B}^t]^T$.

The regular BAMP iteration is executed independently for each of the N_B blocks (jointly sparse signals). Once iteration t is finished for all blocks, a **collective binary extrinsic group update**, similar to (36), is executed to exploit the joint support structure among the N_B signals:

$$\begin{aligned} \bar{L}_{n,b}^t &= U_G(\mathbf{U}^{t-1}, \boldsymbol{\beta}^{t-1}, \gamma_{n,b}^0) := \bar{L}_{n,b}^0 + \sum_{l \in \mathcal{B} \setminus b} L_{n,l}^t \\ &= \log \frac{\gamma_{n,b}^0}{1 - \gamma_{n,b}^0} + \sum_{l \in \mathcal{B} \setminus b} \log \frac{1 - 2u_{n,l}^{t-1}}{2\beta_l^{t-1}}, \end{aligned} \quad \forall n \in \mathcal{V}, \forall b \in \mathcal{B}. \quad (49)$$

Afterwards, the zero probabilities are updated for the subsequent iteration by executing prior update (37), which is now applied entry-wise on a matrix.

The BOSSAMP algorithm for the jointly sparse case is depicted by Algorithm 4.

B. BOSSAMP and Jointly Sparse Gaussian Signals

Considering the block structure indexing, the **collective Gaussian extrinsic group update**, similar to (45), reads

$$\begin{aligned} \bar{L}_{n,b}^t &= U_G(\mathbf{U}^{t-1}, \boldsymbol{\beta}^{t-1}, \gamma_{n,b}^0) := \bar{L}_{n,b}^0 + \sum_{l \in \mathcal{B} \setminus b} L_{n,l}^t \\ &= \log \frac{\gamma_{n,b}^0}{1 - \gamma_{n,b}^0} + \frac{1}{2} \sum_{l \in \mathcal{B} \setminus b} \log \frac{\beta_l^{t-1} + \sigma_{x_{n,l}}^2}{\beta_l^{t-1}} - \frac{(u_{n,l}^{t-1})^2 \sigma_{x_{n,l}}^2}{\beta_l^{t-1} (\beta_l^{t-1} + \sigma_{x_{n,l}}^2)}, \end{aligned} \quad \forall n \in \mathcal{V}, \forall b \in \mathcal{B}. \quad (50)$$

The resulting $\bar{\mathbf{L}}^t$ is then used for the subsequent entry-wise prior update (37).

V. RECOVERY OF ARBITRARY STRUCTURED SIGNALS

While the previous sections presented BOSSAMP for the prominent examples of group/jointly sparse binary and group/jointly sparse Gaussian signals, this section will discuss the generalization to arbitrary signals.

A. Arbitrary Sparse Signals

BOSSAMP is potentially⁵ applicable to arbitrarily distributed signals — the crux is the utilization of the latent variable $z_{n,i}$ as demonstrated in the Gaussian case, see Section III-B. Soft information in terms of L -values is computed and exchanged extrinsically among the group (or joint) structure of the signal(s). In the following, we consider the group sparse case; the jointly sparse case just differs in the indexing as discussed in Section IV-A.

⁵At this point, we can not guarantee the stability of loopy belief propagation and BOSSAMP for all arbitrary prior distributions.

Consider an arbitrary sparse signal prior (similar to (20) and (24))

$$f_{\mathbf{x}_n}(x_n) = \gamma_n \delta(x_n) + (1 - \gamma_n) f_{a_n}(x_n), \quad (51)$$

where $f_{a_n}(x_n)$ is the distribution of the nonzero entries in \mathbf{x} . Remember that random variable u_n is the sum of the two independent random variables \mathbf{x}_n and $\tilde{\mathbf{w}}_n$. The prior distribution of the decoupled measurements (cf. (38) for Gaussian case) is, therefore, obtained via convolution

$$\begin{aligned} f_{u_n}(u_n) &= \int_{\mathbb{R}} f_{\mathbf{x}_n}(v) f_{\tilde{\mathbf{w}}}(u_n - v) dv \\ &= \gamma_n \mathcal{N}(u_n | 0, \beta) + (1 - \gamma_n) f_{a_n * \tilde{\mathbf{w}}_n}(u_n), \end{aligned} \quad (52)$$

where $f_{a_n * \tilde{\mathbf{w}}_n}(u_n) = \int_{\mathbb{R}} f_{a_n}(v) f_{\tilde{\mathbf{w}}}(u_n - v) dv$. Our zero probability estimates (cf. (42) for Gaussian case) now compute as

$$\tilde{\gamma}_n^t = \frac{\gamma_n^{t-1} \mathcal{N}(u_n^{t-1} | 0, \beta^{t-1})}{\gamma_n^{t-1} \mathcal{N}(u_n^{t-1} | 0, \beta^{t-1}) + (1 - \gamma_n^{t-1}) f_{a_n * \tilde{\mathbf{w}}_n}(u_n^{t-1})}, \quad (53)$$

and the innovation L -values are obtained as

$$L_n^t = \log \frac{\tilde{\gamma}_n^t (1 - \gamma_n^{t-1})}{\gamma_n^{t-1} (1 - \tilde{\gamma}_n^t)}. \quad (54)$$

The extrinsic group update is performed with these L -values:

$$\begin{aligned} \bar{L}_n^t &= U_G(\mathbf{u}^{t-1}, \beta^{t-1}, \gamma_n^0) := \bar{L}_n^0 + \sum_{l \in \mathcal{G}_g \setminus n} L_l^t, \\ \forall g \in \{1, \dots, N_G\}. \end{aligned} \quad (55)$$

Afterwards, prior update (37) is executed.

Following this approach using, e.g., the sparse binary prior (20), the resulting L -values (54) coincide with (35). Note, however, that not every arbitrary distribution will entail good recovery performance in the group sparse case; the employed loopy belief propagation at the heart of BOSSAMP may become unstable and require extensions such as damping, see [41], [42].

B. Arbitrarily Structured Signals

Up to now, we considered either group sparse or jointly sparse signals. It is straightforward to extend BOSSAMP to be applicable to jointly sparse signals that, individually, exhibit a group structure. To that end, the group update has to be extended as follows:

$$\begin{aligned} \bar{L}_{n,b}^t &= U_G(\mathbf{U}^{t-1}, \beta^{t-1}, \gamma_{n,b}^0) := \bar{L}_{n,b}^0 + \sum_{i \in \mathcal{G}_g \setminus n} \sum_{j \in \mathcal{B} \setminus b} L_{i,j}^t, \\ \forall g \in \{1, \dots, N_G\}, \forall b \in \mathcal{B}. \end{aligned} \quad (56)$$

As the $|\mathcal{B}| = N_B$ signals are jointly sparse, all of them exhibit the same individual group structure, i.e., the same groups \mathcal{G}_g . Update (56) accounts for the group as well as the joint sparsity.

VI. COMPARISON AND NUMERICAL RESULTS

Let us first introduce the figures of merit for comparison. We then highlight the schemes to which we compare BOSSAMP to. The numerical setup is described, and the simulation results are provided, followed by a discussion. Note that for the sake of brevity, we only present results for the *group sparse* case.

A. Figures of Merit

The measurement Signal-to-Noise Ratio (SNR) is defined as

$$\text{SNR} = \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\mathbb{E}_{\mathbf{w}} \left\{ \|\mathbf{w}\|_2^2 \right\}} = \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{M\sigma_w^2}, \quad (57)$$

the noise variance σ_w^2 is set accordingly for each realization of \mathbf{A} and \mathbf{x} to realize a certain SNR during simulation.

The Normalized Mean Squared Error (NMSE) between original signal \mathbf{x} and its estimate (recovery) $\hat{\mathbf{x}}$ is defined as

$$\text{NMSE} = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2}{\|\mathbf{x}\|_2^2}, \quad (58)$$

it gives indication about the overall recovery performance.

The False Alarm NMSE (FANMSE) is defined as

$$\text{FANMSE} = \frac{\|\mathbf{x}_{\bar{\mathcal{S}}_x} - \hat{\mathbf{x}}_{\bar{\mathcal{S}}_x}\|_2^2}{\|\mathbf{x}\|_2^2} = \frac{\|\hat{\mathbf{x}}_{\bar{\mathcal{S}}_x}\|_2^2}{\|\mathbf{x}\|_2^2}, \quad (59)$$

where the complementary signal support $\bar{\mathcal{S}}_x$ contains the indices of the zero entries in \mathbf{x} . The FANMSE is a measure to quantify the strength of the false alarms in $\hat{\mathbf{x}}$.

B. Comparative Schemes

We compare our implementations of AMP, BAMP and BOSSAMP (MATLAB code will be made available at [43]) to the following schemes:

Group LASSO (GLASSO): in order to incorporate the group structure of \mathbf{x} , the group LASSO [15]–[17] replaces the ℓ_1 -norm regularization in (11) with the sum of ℓ_2 -norms of the groups:

$$\hat{\mathbf{x}}_{\text{GLASSO}}(\mathbf{y}; \lambda) = \arg \min_{\tilde{\mathbf{x}}} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{A}\tilde{\mathbf{x}}\|_2^2 + \lambda \sum_{g=1}^{N_G} \|\tilde{\mathbf{x}}_{\mathcal{G}_g}\|_2 \right\}. \quad (60)$$

In case of $|\mathcal{G}_g| = 1$ and $N_G = N$, it collapses to the standard LASSO (11).

We use an implementation via the alternating direction method of multipliers that is described in [17] and whose MATLAB code is available in [44].

A similar approach to solve the group LASSO via the alternating direction method was presented in [18]; however, our simulations have shown that [44] yields superior results.

Hybrid Generalized Approx. Message Passing (HGAMP): Generalized Approximate Message Passing (GAMP) was introduced in [11] to extend the classical Gaussian AMP framework – on which we build in this paper – to a more general setting that allows for arbitrary output channels and is thus not restricted to additive Gaussian noise in (2). An extension to incorporate structured sparsity was introduced in [20] and is termed HGAMP. It was shown to outperform group orthogonal matching pursuit [19] and the group LASSO in terms of NMSE (58). We use the MATLAB implementation of HGAMP that is provided in [45].

C. Numerical Setup

For AMP, BAMP and BOSSAMP, the stopping criterion was set to $\epsilon_{\text{tol}} = 10^{-4}$. The maximum number of iterations was set to $t_{\text{max}} = 100$, for all algorithms.

For AMP Algorithm 1, we chose $\lambda = 2.678K^{-0.181}$ (NMSE minimizing heuristic for $N = 1000$, see [46]).

For GLASSO, the regularization parameter λ is chosen according to the example provided in [44], and the augmented Lagrangian and over-relaxation parameters were chosen as $\rho = 1$ and $\alpha = 1$, respectively, as suggested in [44].

For HGAMP with *sparse Gaussian* signal prior, following options were selected (suggested by the toy example in [45]): `step=1`, `removeMean=true`, `adaptStep=true`. In case of *sparse binary* signal prior, following options were changed in order to mitigate numerical issues: `step=0.1`, `removeMean=false`. In function `estim` of class `GrpSparseEstim.m`, the minimum and maximum value of the sparse probability `rho` was set to 10^{-12} and $1 - 10^{-12}$, respectively; the same values were chosen for the minimum and maximum value of `pr0` — this improved the recovery performance, and mitigated numerical issues in the binary case.

For the “**variable SNR**” and “**variable M** ” curves, the results are averaged over 1000 random realizations. In each realization, sensing matrix \mathbf{A} and signal vector \mathbf{x} are newly generated: \mathbf{A} features i.i.d. zero mean Gaussian entries with unit ℓ_2 -norm columns, and \mathbf{x} has dimension $N = 1000$ with $K = 160$ nonzero entries, entailing a zero probability of $\gamma_n = 0.84$. The nonzero entries are one in the *sparse binary* case (20), and i.i.d. Gaussian with zero mean and variance $\sigma_{x_n}^2 = 1$ in the *sparse Gaussian* case (24). We consider non-overlapping equally-sized groups in \mathbf{x} and compare three different cases:

- Group size $|\mathcal{G}_g| = 2, \forall g \in \{1, 2, \dots, 500\}$. With $K = 160$, this implies that we have 80 active groups out of $N_G = 500$ total groups.
- Group size $|\mathcal{G}_g| = 5, \forall g \in \{1, \dots, 200\}$ (32 active groups).
- Group size $|\mathcal{G}_g| = 8, \forall g \in \{1, \dots, 125\}$ (20 active groups).

For the “**empirical phase transition**” curves, we consider an undersampling $(\frac{M}{N})$ vs. sparsity $(\frac{K}{N})$ grid, where the values range from 0.05 to 0.95 with stepsize 0.05, respectively. At each grid point, 200 realizations are simulated. Let us introduce a *success indicator* for each realization r :

$$S_r = \begin{cases} 1 & \text{NMSE}_r < 10^{-4} \\ 0 & \text{else} \end{cases}. \quad (61)$$

The *average success* is obtained as $\bar{S} = \frac{1}{200} \sum_{r=1}^{200} S_r$. The empirical phase transition curves are finally obtained by plotting the 0.5 contour of \bar{S} using MATLAB function `contour`.

D. Numerical Results and Discussion

For convenience, Table I highlights how the various schemes utilize prior information, i.e., the sparsity, the Bayesian prior (see (20) and (24)), and the group structure. In the following, we call BAMP, BOSSAMP and GAMP the Bayesian message passing-based schemes.

	Sparsity	Bayesian prior	Group structure
GLASSO	✓	✗	✓
AMP	✓	✗	✗
BAMP	✓	✓	✗
BOSSAMP	✓	✓	✓
HGAMP	✓	✓	✓

TABLE I
UTILIZATION OF PRIOR INFORMATION.

Figure 4 shows the **variable SNR** results for the *sparse binary* case, where the number of measurements was fixed to $M = 590$ (inspired by (4) with $c \approx 2$, see [46]). BOSSAMP, HGAMP and GLASSO depend on the group size that is indicated in brackets, while AMP and BAMP do not exploit the group structure. AMP and GLASSO only include a sparsity constraint steered by λ in (11) and (60), respectively. BAMP, BOSSAMP and HGAMP exploit the full prior knowledge, i.e., they utilize prior (20). Doing so, these schemes exhibit a steep transition to the success state ($\text{NMSE} < -40$ dB), provided that the SNR is sufficiently high. It is evident that larger groups strongly improve the results of the schemes that exploit the group structure, i.e., BOSSAMP, HGAMP and GLASSO. The FANMSE draws a similar picture as the NMSE. It is notable that the GLASSO is able to effectively null the false alarms, given reasonably sized groups. Once the SNR is large enough to ensure successful recovery, the number of iterations of the Bayesian message passing-based schemes stays very low. The overall best recovery performance is obtained by our proposed BOSSAMP algorithm, followed by HGAMP and BAMP.

Figure 5 shows the **variable M** results for the *sparse binary* case at $\text{SNR} = 25$ dB. We observe very steep success transitions for the Bayesian message passing-based schemes. In particular, BOSSAMP with group size 2 yields successful recoveries above 140 measurements, while for group size 8, it only requires slightly more than 20 measurements. In comparison, HGAMP is successful above 200 and 40 measurements, respectively, and BAMP requires around 300 measurements.

Figure 6 shows the **variable SNR** results for the *sparse Gaussian* case with $M = 590$. The message passing-based schemes do not exhibit the same steep success transitions as in the sparse binary case in Figure 4, but a gradually decreasing NMSE over increasing SNR. The performance of BOSSAMP and HGAMP is very similar, particularly for large group sizes. The FANMSE shows a steeper decrease over SNR, and it is again notable that GLASSO features a steep decay, at the expense of stagnating NMSE (the algorithm utilizes thresholding that leads to a sparse solution but lowers the energy in the nonzero entries). GLASSO overall requires the least number of iterations, closely followed by BOSSAMP. However, the Bayesian message passing-based schemes strongly outperform AMP and GLASSO in terms of NMSE and FANMSE.

Figure 7 depicts the **variable M** results for the *sparse Gaussian* case at $\text{SNR} = 25$ dB. The steep success transitions of the Bayesian message passing-based schemes are back, the




Fig. 4. Variable SNR in the sparse binary case.

NMSE is lower bounded due to finite SNR. It is again evident that BOSSAMP features earlier phase transitions (requires fewer measurements) than HGAMP, which is particularly apparent at small group sizes. The number of iterations behave similarly for BOSSAMP, HGAMP and BAMP; a distinctive peak accompanies the phase transition event, after which the iterations decrease. Due to the same message passing foundation, BOSSAMP and HGAMP behave similarly.

Figure 8 illustrates the **empirical phase transition** curves for the *sparse binary* case. The standard AMP algorithm exhibits the worst performance and is strongly surpassed by BAMP that incorporates the Bayesian prior knowledge. Additionally exploiting the group structure leads to a supreme

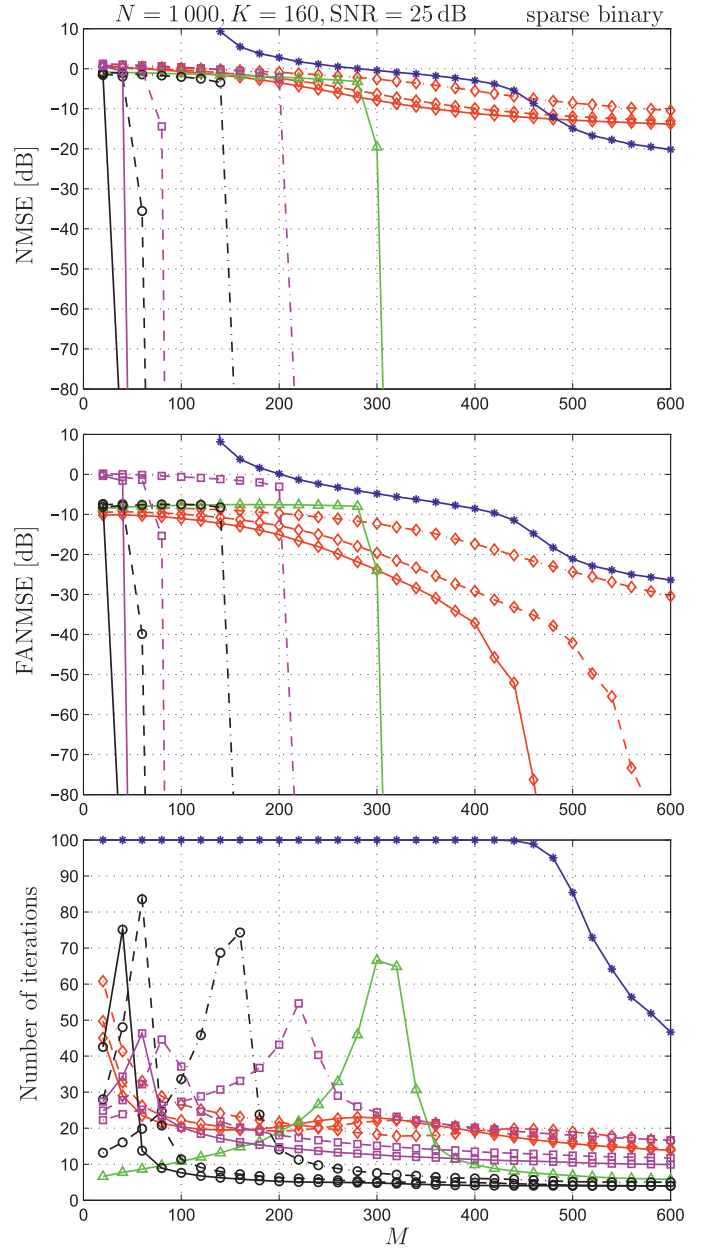


Fig. 5. Variable M in the sparse binary case.

performance which can be seen in the BOSSAMP phase transition curve. Note that the group size in this example was chosen really small as $|\mathcal{G}_g| = 2$, yet already results in a big improvement.

Finally, Figure 9 illustrates the **empirical phase transition** curves for the *sparse Gaussian* case for various group sizes. Clearly, an increase in the group size strongly improves the recovery performance of BOSSAMP and HGAMP and leads to earlier success transitions. While BOSSAMP and HGAMP behave similarly at large values $\frac{M}{N}$, the strongly undersampled regime causes problems for HGAMP.




Fig. 6. Variable SNR in the sparse Gaussian case.

VII. CONCLUSION

We introduced BOSSAMP, a novel iterative algorithm to efficiently recover group sparse or jointly sparse signals. The algorithm is based on the AMP framework introduced by Donoho, Maleki and Montanari and exploits the known signal prior distribution. By introducing an *extrinsic group update* and a *prior update* step in each iteration, the known signal structure is incorporated into the entry-wise MMSE estimation of the standard BAMP algorithm; considering a specific element of a group, the extrinsic group update step collects soft information from the remaining group elements. L -values are accumulated according to the turbo principle and a belief about whether a specific group element was zero




Fig. 7. Variable M in the sparse Gaussian case.

or nonzero arises. According to this belief, the subsequent prior update step updates the zero probability of the prior distribution that is utilized for the MMSE estimation in BAMP.

We derived the group update step for the sparse binary respectively the sparse Gaussian case and provided simple closed form expressions. Furthermore, we sketched how BOSSAMP is potentially applicable to arbitrary sparse signals. Simulations have shown that BOSSAMP outperforms current state of the art algorithms, including HGAMP that builds on the same message passing foundation. However, HGAMP is based on GAMP that – in contrast to the standard AMP framework that we utilize in this work – is applicable not only to the additive Gaussian noise case but to a more general class of problems,




Fig. 8. Empirical phase transition curves for sparse binary prior. Recoveries are successful ($\text{NMSE} < 10^{-4}$) in the regime below a curve.




Fig. 9. Empirical phase transition curves for sparse Gaussian prior. Recoveries are successful ($\text{NMSE} < 10^{-4}$) in the regime below a curve. Results are plotted for three different group sizes $|G_g| = \{2, 5, 8\}$.

including nonlinear output relations. While being more general, HGAMP encompasses a more difficult implementation, and simulations suggest that due to a series of (additional) approximations, it sacrifices some performance in comparison to the standard AMP framework [6]–[9].

Currently, the signal prior distribution is assumed to be known. For the sparse Gaussian and the more general Gaussian mixture case, the parameters can be estimated using the EM-algorithm, whose application in conjunction with GAMP has been propagated in [47], [48]. Another promising and more general approach was proposed in [49], where Stein’s unbiased risk estimator was incorporated into BAMP.

In conclusion, we have shown that the utilization of the (known) signal structure leads to significant improvements — on the one hand, fewer measurements are required to obtain a certain recovery performance (improved phase transition), while on the other hand, the recovery becomes more robust with respect to noise if the number of measurements is fixed. BOSSAMP is a versatile, easy-to-implement recovery algorithm with great performance.

REFERENCES

- [1] D. L. Donoho, “Compressed Sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [2] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [3] E. J. Candès, J. K. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on pure and applied mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [4] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic Decomposition by Basis Pursuit,” *SIAM journal on scientific computing*, vol. 20, no. 1, pp. 33–61, 1998.
- [5] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- [6] D. L. Donoho, A. Maleki, and A. Montanari, “Message-passing algorithms for compressed sensing,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18914–18919, 2009.
- [7] M. A. Maleki, *Approximate message passing algorithms for compressed sensing*. PhD Thesis, Stanford University, 2010.
- [8] D. L. Donoho, A. Maleki, and A. Montanari, “Message Passing Algorithms for Compressed Sensing: I. Motivation and Construction,” in *2010 IEEE Information Theory Workshop (ITW)*, pp. 1–5.
- [9] D. L. Donoho, A. Maleki, and A. Montanari, “How to Design Message Passing Algorithms for Compressed Sensing,” *preprint*, 2011.
- [10] Y. Weiss and W. T. Freeman, “Correctness of belief propagation in Gaussian graphical models of arbitrary topology,” *Neural computation*, vol. 13, no. 10, pp. 2173–2200, 2001.
- [11] S. Rangan, “Generalized approximate message passing for estimation with random linear mixing,” in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 2168–2172, IEEE, 2011.
- [12] S. Zhang, J. Huang, Y. Huang, Y. Yu, H. Li, and D. N. Metaxas, “Automatic image annotation using group sparsity,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3312–3319, IEEE, 2010.
- [13] B. Ng and R. Abugharbieh, “Generalized group sparse classifiers with application in fMRI brain decoding,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1065–1071, IEEE, 2011.
- [14] J. Gui, D. Tao, Z. Sun, Y. Luo, X. You, and Y. Y. Tang, “Group sparse multiview patch alignment framework with view consistency for image classification,” *IEEE Transactions on Image Processing*, vol. 23, no. 7, pp. 3126–3137, 2014.
- [15] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [16] J. Friedman, T. Hastie, and R. Tibshirani, “A note on the group lasso and a sparse group lasso,” *arXiv preprint arXiv:1001.0736*, 2010.
- [17] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [18] W. Deng, W. Yin, and Y. Zhang, “Group sparse optimization by alternating direction method,” in *SPIE Optical Engineering+ Applications*, pp. 88580R–88580R, International Society for Optics and Photonics, 2013.
- [19] G. Swirszcz, N. Abe, and A. C. Lozano, “Grouped orthogonal matching pursuit for variable selection and prediction,” in *Advances in Neural Information Processing Systems*, pp. 1150–1158, 2009.
- [20] S. Rangan, A. K. Fletcher, V. K. Goyal, and P. Schniter, “Hybrid generalized approximate message passing with applications to structured sparsity,” in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 1236–1240, IEEE, 2012.
- [21] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado, “Sparse solutions to linear inverse problems with multiple measurement vectors,” *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2477–2488, 2005.
- [22] J. Ziniel and P. Schniter, “Efficient high-dimensional inference in the multiple measurement vector problem,” *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 340–354, 2013.
- [23] D. Liang, L. Ying, and F. Liang, “Parallel MRI Acceleration Using M-FOCUSS,” in *3rd International Conference on Bioinformatics and Biomedical Engineering (ICBBE)*, pp. 1–4, IEEE, 2009.

- [24] G. Tzagkarakis, D. Miliotis, and P. Tsak, "Multiple-measurement Bayesian compressed sensing using GSM priors for DOA estimation," in *IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, pp. 2610–2613, IEEE, 2010.
- [25] J. Hagenauer, "Source-controlled channel decoding," *IEEE Transactions on Communications*, vol. 43, no. 9, pp. 2449–2457, 1995.
- [26] J. Hagenauer, E. Offer, and L. Papke, "Iterative decoding of binary block and convolutional codes," *IEEE Transactions on Information Theory*, vol. 42, no. 2, pp. 429–445, 1996.
- [27] J. Hagenauer, "The exit chart-introduction to extrinsic information transfer in iterative processing," in *Proc. 12th European Signal Processing Conference (EUSIPCO)*, pp. 1541–1548, Citeseer, 2004.
- [28] C. Berrou and A. Glavieux, "Near optimum error correcting coding and decoding: Turbo-codes," *Communications, IEEE Transactions on*, vol. 44, no. 10, pp. 1261–1271, 1996.
- [29] P. Schniter, "Turbo reconstruction of structured sparse signals," in *44th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6, IEEE, 2010.
- [30] S. Som and P. Schniter, "Compressive imaging using approximate message passing and a markov-tree prior," *IEEE Transactions on Signal Processing*, vol. 60, no. 7, pp. 3439–3448, 2012.
- [31] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
- [32] A. Montanari, "Graphical Models Concepts in Compressed Sensing," *Compressed Sensing: Theory and Applications*, pp. 394–438, 2012.
- [33] J. S. Yedidia, "Message-passing algorithms for inference and optimization," *Journal of Statistical Physics*, vol. 145, no. 4, pp. 860–890, 2011.
- [34] H.-A. Loeliger, J. Dauwels, J. Hu, S. Korl, L. Ping, and F. R. Kschischang, "The Factor Graph Approach to Model-Based Signal Processing," *Proceedings of the IEEE*, vol. 95, no. 6, pp. 1295–1322, 2007.
- [35] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 498–519, 2001.
- [36] A. Mousavi, A. Maleki, and R. G. Baraniuk, "Asymptotic Analysis of LASSOs Solution Path with Implications for Approximate Message Passing," *arXiv preprint arXiv:1309.5979*, 2013.
- [37] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Transactions on Information Theory*, vol. 57, no. 2, pp. 764–785, 2011.
- [38] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the royal statistical society. Series B (methodological)*, pp. 1–38, 1977.
- [39] C. M. Bishop, *Pattern recognition and machine learning*, vol. 4. springer New York, 2006.
- [40] L. Xu and M. I. Jordan, "On convergence properties of the EM algorithm for Gaussian mixtures," *Neural computation*, vol. 8, no. 1, pp. 129–151, 1996.
- [41] T. Heskes, "Stable fixed points of loopy belief propagation are local minima of the bethe free energy," in *Advances in neural information processing systems*, pp. 343–350, 2002.
- [42] T. Heskes, "On the uniqueness of loopy belief propagation fixed points," *Neural Computation*, vol. 16, no. 11, pp. 2379–2413, 2004.
- [43] M. Mayer and N. Goertz, *MATLAB code for AMP, BAMP and BOSSAMP*. <https://www.nt.tuwien.ac.at/downloads/featured-downloads>.
- [44] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *ADMM group LASSO implementation*. <https://web.stanford.edu/~boyd/papers/admm/>.
- [45] S. Rangan, A. Fletcher, V. Goyal, U. Kamilo, J. Parker, P. Schniter, J. Vila, J. Ziniel, and M. Borgering, *GAMP Implementation*. <http://gampmatlab.sourceforge.net/>.
- [46] M. Mayer, N. Goertz, and J. Kaitovic, "RFID Tag Acquisition via Compressed Sensing," in *Proceedings of IEEE RFID Technology and Applications Conference (RFID-TA)*, pp. 26–31, 2014.
- [47] J. Vila and P. Schniter, "Expectation-maximization Bernoulli-Gaussian approximate message passing," in *Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, pp. 799–803, IEEE, 2011.
- [48] J. P. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," *IEEE Transactions on Signal Processing*, vol. 61, no. 19, pp. 4658–4672, 2013.
- [49] C. Guo and M. Davies, "Near Optimal Compressed Sensing Without Priors: Parametric SURE Approximate Message Passing," *IEEE Transactions on Signal Processing*, vol. 63, pp. 2130–2141, April 2015.